1. An integrated circuit functioning as an InfiniBand channel adapter and an InfiniBand switch, comprising:

a plurality of InfiniBand media access controllers (MACs), for transceiving InfiniBand packets;

a plurality of local bus interfaces, for performing addressed data transfers on a plurality of local buses coupled thereto;

a bus router, for performing transport layer operations between said plurality of InfiniBand MACs and said plurality of local bus interfaces; and

a transaction switch, coupled to each of said plurality of InfiniBand MACs, said plurality of local bus interfaces, and said bus router, for switching data and transactions therebetween.

2. The integrated circuit of claim 1, further comprising:

a plurality of transaction queues, associated with said plurality of InfiniBand MACs, said plurality of local bus interfaces, and said bus router, coupled to said transaction switch, for storing said transactions.

3. The integrated circuit of claim 1, further comprising:

a memory, shared by said plurality of InfiniBand MACs, said plurality of local bus interfaces, and said bus router, for buffering data received thereby.

4. The integrated circuit of claim 3, wherein said transaction switch comprises:

a buffer manager, for allocating portions of said memory to said plurality of InfiniBand MACs, said plurality of local bus interfaces, and said bus router, for buffering said data received thereby.

5. The integrated circuit of claim 4, wherein said buffer manager performs said allocating in a substantially as-needed manner.

6. The integrated circuit of claim 3, wherein said bus router is configured to write an InfiniBand packet header into said memory via said transaction switch along addressed data stored in said memory by one of said plurality of local bus interfaces to create an InfiniBand packet.

7. The integrated circuit of claim 3, wherein said plurality of local bus interfaces are configured to read a payload portion of an InfiniBand packet stored in said memory and to transmit said payload portion on one or more of the plurality of local buses coupled thereto.

8. The integrated circuit of claim 7, wherein said payload portion is located in said memory at an offset

specified in a transaction posted by said bus router to said plurality of local bus interfaces via said transaction switch.

9. The integrated circuit of claim 1, wherein at least one of said plurality of local bus interfaces comprises a PCI bus interface.

10. The integrated circuit of claim 1, wherein said transaction switch is configured to receive a transaction posted by a first of said plurality of InfiniBand MACs in response to a packet received by said first of said plurality of InfiniBand MACs and to selectively switch said transaction to one of a second of said plurality of InfiniBand MACs and said bus router.

11. The integrated circuit of claim 10, wherein said transaction switch selectively switches said transaction based on an InfiniBand destination local identification value included in said transaction.

12. The integrated circuit of claim 11, wherein said transaction switch selectively switches said transaction to said bus router if an entry associated with said InfiniBand destination local identification value in a mapping table of said transaction switch

indicates said transaction is destined for said bus router.

13. The integrated circuit of claim 12, wherein said transaction switch selectively switches said transaction to one of said plurality of InfiniBand MACs based on which of said plurality of InfiniBand MACs is associated with said InfiniBand destination local identification value in said mapping table if said entry indicates said transaction is not destined for said bus router.

14. The integrated circuit of claim 11, wherein said first MAC parses said InfiniBand destination local identification value from said packet.

15. The integrated circuit of claim 10, wherein said transaction includes an InfiniBand virtual lane number parsed from said packet.

16. The integrated circuit of claim 10, wherein said transaction includes a destination queue pair number parsed from said packet.

17. The integrated circuit of claim 1, wherein said transaction switch is configured to receive a transaction posted by said bus router and to selectively switch said transaction to one of said

plurality of InfiniBand MACs and one of one of said plurality of local bus interfaces.

18. The integrated circuit of claim 17, wherein said transaction switch selectively switches said transaction based on a transaction type value included in said transaction.

19. The integrated circuit of claim 18, wherein said transaction switch selectively switches said transaction to one of said plurality of local bus interfaces based on whether a local bus address included in said transaction falls into one or more predetermined address ranges of the plurality of local buses.

20. The integrated circuit of claim 17, wherein said transaction includes an address in an address range of the plurality of local buses.

21. The integrated circuit of claim 1, further comprising a local bus bridge coupled between said plurality of local bus interfaces for buffering data therebetween.

22. The integrated circuit of claim 1, wherein said transaction switch is configured to receive a transaction posted by a first of said plurality of local bus interfaces in response to an addressed data

transfer received by said first of said plurality of local bus interfaces and to switch said transaction to a second of said plurality of local bus interfaces.

23. A transaction switch for switching data between a plurality of data devices, comprising:

a memory, shared by the plurality of data devices for buffering data received thereby;

multiplexing logic, for controlling the transfer of data between the plurality of data devices and said memory; and

control logic, for controlling said multiplexing logic;

wherein the plurality of data devices comprise a plurality of packetized data devices and a plurality of addressed data devices;

wherein said control logic is configured to selectively control said multiplexing logic to transfer data through said memory between two of said packetized data devices and between one of said packetized data devices and one of said addressed data devices.

24. The transaction switch of claim 23, wherein said control logic is further configured to selectively control said multiplexing logic to transfer data

through said memory between two of said addressed data devices.

25. The transaction switch of claim 23, wherein said control logic is configured to selectively control said multiplexing logic to transfer data through said memory between two of said packetized data devices and between one of said packetized data devices and one of said addressed data devices concurrently.

26. The transaction switch of claim 23, wherein at least two of said packetized data devices comprise InfiniBand interfaces.

27. The transaction switch of claim 23, wherein at least two of said addressed data devices comprise PCI bus interfaces.

28. The transaction switch of claim 23, further comprising:

a buffer manager, for allocating portions of said memory to the plurality of data devices for buffering said data.

29. The transaction switch of claim 28, wherein said buffer manager is configured to perform said allocating on substantially a first-come-first-serve basis.

30. The transaction switch of claim 23, wherein said control logic is configured to selectively control said multiplexing logic to transfer data through said memory between two of said packetized data devices and between one of said packetized data devices and one of said addressed data devices in response to a transaction posted to the transaction switch by the plurality of data devices.

31. The transaction switch of claim 30, wherein said transaction comprises a command to transfer data between said memory and one of the plurality of data devices.

32. The transaction switch of claim 31, wherein said transaction comprises an address of a buffer within said memory wherein is stored said data to be transferred in response to said command.

33. The transaction switch of claim 32, wherein said transaction comprises an offset within said buffer for addressing portions of said data.

34. The transaction switch of claim 30, wherein said transaction comprises a tag for uniquely identifying said transaction from other transactions posted to the transaction switch by the plurality of data devices.

35. The transaction switch of claim 23, wherein the plurality of data devices comprise a transport layer device, wherein the transaction switch is configured to receive transactions from said transport layer device for performing protocol translation of data between said one of said packetized data devices and said one of said addressed data devices.

36. A transaction switch for switching transactions and data between a plurality of data interfaces, the transaction switch comprising:

a memory, shared by the plurality of data interfaces, for buffering data received thereby;

a plurality of transaction queues, associated with each of the plurality of data interfaces, configured to store transactions, said transactions adapted to convey information to enable the plurality of data interfaces to transfer said data according to a plurality of disparate data transfer protocols supported thereby; and

control logic, configured to route said data through said shared memory between the plurality of data interfaces and to switch said transactions between the plurality of data interfaces.

37. The transaction switch of claim 36, wherein said control logic is configured to route said data between the plurality of data interfaces through said shared memory in response to said transactions received from said plurality of transaction queues.

38. The transaction switch of claim 36, wherein at least a portion of said plurality of transaction queues is configured to store transactions adapted to convey information necessary to transfer data according to an InfiniBand protocol.

39. The transaction switch of claim 36, wherein at least a portion of said plurality of transaction queues is configured to store transactions adapted to convey information necessary to transfer data according to an PCI bus protocol.

40. The transaction switch of claim 36, wherein said control logic is further configured to modify a transaction received from one of said plurality of transaction queues associated with a first of the plurality of data devices and to send said modified transaction to another one of said plurality of transaction queues.

41. An integrated circuit, comprising:

at least three data interfaces;

a memory, shared by said at least three data interfaces for buffering data therebetween; and

a transaction switch, coupled to said at least three data interfaces and said memory, for dynamically allocating portions of said memory to said at least three data interfaces for storing data therein, and for controlling access to said allocated portions of said memory by each of said at least three data interfaces;

wherein at least one of said at least three data interfaces is of a different type than the others.

42. The integrated circuit of claim 41, wherein at least one of said at least three data interfaces is a packetized data interface and at least one of said at least three data interfaces is an addressed data interface.

43. The integrated circuit of claim 42, wherein said at least one packetized data interface is an InfiniBand interface.

44. The integrated circuit of claim 42, wherein said at least one addressed data interface is a PCI interface.

45. The integrated circuit of claim 41, wherein said transaction switch is configured to receive a transaction from a first of said at least three data interfaces and to selectively switch said transaction to one of another of said at least three data interfaces.

46. The integrated circuit of claim 45, wherein said transaction is a packetized data transaction including packet destination information.

47. The integrated circuit of claim 46, wherein said transaction switch is configured to selectively switch said packetized data transaction to said another of said at least three data interfaces based on said packet destination information and information stored in a mapping table of said transaction switch.

48. The integrated circuit of claim 47, wherein said transaction switch is configured to selectively switch said packetized data transaction to said another of said at least three data interfaces further based on information stored in a table mapping said packet destination information to said at least three data interfaces.

49. The integrated circuit of claim 45, wherein in a first instance of said transaction said first and one of another of said at least three data interfaces are of a same type of interface, wherein in a second instance of said transaction said first and one of another of said at least three data interfaces are of a different type of interface.

50. The integrated circuit of claim 49, wherein in said first instance each of said first and one of another of said at least three data interfaces is a packetized data interface type.

51. The integrated circuit of claim 49, wherein in said second instance said first of said at least three data interfaces is a packetized data interface type and said one of another of said at least three data interfaces is an interface type capable of translating between packetized and addressed data.

52. The integrated circuit of claim 49, wherein in said second instance said first of said at least three data interfaces is a packetized data interface type and said one of another of said at least three data interfaces is a transport level data interface.

53. The integrated circuit of claim 45, wherein said transaction switch is configured to modify said transaction received from said first of said at least three data interfaces prior to selectively switching said received transaction to said one of another of said at least three data interfaces.

54. The integrated circuit of claim 41, wherein said transaction switch is configured to receive a transaction from a first of said at least three data interfaces and to selectively switch said received transaction to two or more of another of said at least three data interfaces.

55. The integrated circuit of claim 41, further comprising:
a plurality of transaction queues, coupled between said transaction switch and said at least three data interfaces, for storing transactions between said transaction switch and said at least three data interfaces.

56. The integrated circuit of claim 55, further comprising a programmable register for specifying for at least a plurality of said plurality of transaction queues a number of transaction slots to be allocated for storing said transactions.

57. The integrated circuit of claim 41, wherein at least one of said at least three data interfaces comprises a bus router for performing a transport layer function between at least two other of said at least three data interfaces which support disparate data protocols.

58. The integrated circuit of claim 41, further comprising:
a bus router, coupled to the transaction switch, for performing transport layer functions between said at least three data interfaces having different data protocols.

59. The integrated circuit of claim 58, wherein said bus router is configured to write packet header information into said allocated portions of said memory.

60. The integrated circuit of claim 41, wherein a first of said at least three data interfaces is configured to post a transaction to said transaction switch for instructing a second of said at least three data interfaces to transfer data to or from an offset in one of said allocated portions of said memory associated with a payload portion of a data packet.

61. The integrated circuit of claim 41, wherein said transaction switch is further configured to de-allocate said portions of said memory.

62. The integrated circuit of claim 41, wherein said transaction switch is configured to dynamically allocate said portions of said memory to said at least three data interfaces on a substantially as needed basis.

63. An integrated circuit, comprising:
a plurality of packetized data interfaces;
a plurality of addressed data interfaces; and
a transaction switch, coupled to said plurality of packetized data interfaces and to said plurality of addressed data interfaces;
wherein said transaction switch is configured to switch a packetized data transaction from a first of said plurality of packetized data interfaces to a second of said plurality of packetized data interfaces and to switch an addressed data transaction from a first of said addressed data interfaces to a second of said addressed data interfaces.

64. The integrated circuit of claim 63, further comprising:
a routing device, coupled to said transaction switch, for performing protocol translation between said plurality of packetized data interfaces and said plurality of addressed data interfaces;

wherein said transaction switch is further configured to switch a packetized data transaction from said first of said plurality of packetized data interfaces to said routing device.

65. The integrated circuit of claim 64, wherein said transaction switch is further configured to switch an addressed data transaction from said routing device to said second of said plurality of addressed data interfaces.

66. The integrated circuit of claim 63, wherein each of said plurality of packetized data interfaces comprises a packetized interface selected from a list comprising an Ethernet interface, a FibreChannel interface, an IEEE 1394 interface, a SONET interface, an ATM interface, a SCSI interface, a serial ATA interface, an OC-48 interface, and an OC-192 interface.

67. An InfiniBand hybrid channel adapter/switch, comprising:

a plurality of InfiniBand ports;

at least one addressed data bus interface;

a memory, for buffering data received by said plurality of InfiniBand ports and said at least one addressed data bus interface;

a transport layer engine for routing said data between said plurality of InfiniBand ports and said at least one addressed data bus interface;

a plurality of transaction queues, associated with each of said plurality of InfiniBand ports, said at least one addressed data bus interface and said transport layer engine, for storing transactions; and

a transaction switch, coupled to said plurality of transaction queues, configured to route said transactions between said plurality of InfiniBand ports, said at least one addressed data bus interface and said transport layer engine.

68. The InfiniBand hybrid channel adapter/switch of claim 67, wherein at least a subset of said plurality of transaction queues comprise an input queue for said

plurality of InfiniBand ports to post said transaction to said switch.

~~69. The InfiniBand hybrid channel adapter/switch of claim~~ 67, wherein said plurality of transaction queues comprise an output queue for said transaction switch to send said transaction to said plurality of InfiniBand ports and said at least one addressed data bus interface.

70. The InfiniBand hybrid channel adapter/switch of claim 67, wherein said at least one addressed data bus interface is coupled to a data bus selected from a list comprising a Rapid I/O bus, a VESA bus, an ISA bus, a PCI bus, an LDT bus, an SDRAM bus, a DDR SDRAM bus, and a RAMBUS.

71. A method of transferring data in a network device comprising a plurality of packetized data devices, a plurality of addressed data devices, and at least one routing device, the method comprising:

switching a packetized data transaction from a first of the plurality of packetized data devices to a second of the plurality of packetized data devices; and

switching an addressed data transaction from a first of the plurality of addressed data devices to a second of the plurality of addressed data devices.

72. The method of claim 71, wherein each of said switching steps is performed substantially concurrently.

73. The method of claim 71, wherein said packetized data transaction comprises a transaction associated with an InfiniBand packet transmission.

74. The method of claim 71, wherein said addressed data transaction comprises a transaction associated with a PCI bus data transfer.

75. The method of claim 71, further comprising:

switching a packetized data transaction from one of the plurality of packetized data devices to the routing device.

76. The method of claim 71, further comprising:

switching an addressed data transaction from the routing device to one of the plurality of addressed data devices; and

switching a packetized data transaction from the routing device to one of the plurality of packetized data devices.

77. The method of claim 71, further comprising:

performing protocol translation in a transfer of data between the packetized data interface and the addressed data interface without double-buffering the data.

78. The method of claim 71, further comprising:

parsing a packet and generating said packetized data transaction prior to said switching said packetized data transaction.

79. A transaction switch in a network device having a buffer memory and plurality of data devices, including packetized and addressed data devices, the transaction switch comprising:

a buffer manager, for allocating portions of the buffer memory to the plurality of data devices on an as-needed basis;

a plurality of data paths, for providing the plurality of data devices access to the buffer memory;

a mapping table, for storing packet destination identification information;

a plurality of transaction queues, for transferring transactions between the transaction switch and the plurality of data devices; and

control logic, for selectively switching data between the plurality of data devices based on said mapping table information and in response to said transactions.